

ADAPTIVE INTERPLANETARY NAVIGATION USING GENETIC ALGORITHMS

Todd A. Ely^{*}, Robert H. Bishop[†], and Timothy P. Crain[‡]

The problem of tuning trajectory determination models for interplanetary navigation is a complex task requiring an intensive search of multiple dynamical and nondynamical models that yield trajectory solutions with minimal errors. The process that operational teams currently utilize is based as much on previous experience, as it is on a scientific understanding of these underlying models. This study illustrates an automated approach for filter tuning (via model optimization) using a genetic algorithm (GA) coupled with an extended Kalman filter (EKF). In particular, the solar radiation pressure (SRP) model of the Mars Pathfinder (MPF) spacecraft is investigated using a 3 month span of tracking data during the cruise phase of the mission. The results obtained in this study are compared to the best model obtained by the MPF navigation team. Unlike many previous adaptive filtering schemes, the GA based approach does not require gradient information about neighboring model options, hence it is capable of examining filter models of varying structure. The GA operates on a population of individuals that are selected (initially at random) from the design space. In this study, the selected design space includes $1.44\text{E}+17$ distinct SRP models. Each individual processes the tracking data set using the EKF. The basis for the GA's fitness function is a normalized sample statistic of the output residual sequence. Using the fitness values computed for each individual, the GA selects the parent population via a tournament method. For crossover, several strategies are investigated to determine the best method for quick convergence of the GA to a near optimal solution. The results show that the GA is able to determine an SRP model with a fitness value that is $\sim 6\%$ better than the model selected by the MPF navigation team, and produces predicted residuals that are more stable.

INTRODUCTION

The problem of tuning trajectory determination models for interplanetary navigation is a complex task requiring an intensive search of multiple dynamical and nondynamical models that yield trajectory solutions with minimal errors. The process that operational teams currently utilize is based as much on previous experience, as it is on a scientific understanding of these underlying models. Furthermore, there is an ever-increasing demand for navigation analysts to support multiple spacecraft missions (each with unique modeling issues) using tools that have inherent limitations because of the generic models utilized by these tools. As an example, the Mars Pathfinder (MPF) spacecraft that successfully landed on Mars in July 1997 had a backshell shroud that protected the lander during cruise and Mars entry. It was conical in shape and, during the cruise phase, often shadowed from the Sun by the spacecraft solar arrays. However, the component models available for use by the MPF analysts for solar radiation pressure modeling consisted of only flat plates and cylinders. Neither was entirely correct, thus they were forced to approximate the model of the backshell using one of these choices with an unknown scale factor (because of shadowing) [1]. Their approach was iterative requiring the team to select a model, adjust the filter realization, process the observations, and then compare results to previous filter realizations. Clearly, this tuning process could have benefited from a systematic and automated methodology for finding a best filter model. Doing so would have eased operation team workloads and have allowed them to consider a wider range of possible

^{*} Staff Engineer, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

[†] Associate Professor, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, TX 78712.

[‡] Graduate Student, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, TX 78712.

solutions. This study presents a method for adaptation of navigation models using genetic algorithms to search a selected design space, and arrive at a model that best fits the measured spacecraft tracking data.

Popular approaches to model adaptation often utilize parallel filter banks, each operating with different internal model realizations. The Magill filter bank is a classic method utilizing a Bayesian method to assign probabilities to each member of the bank, with the aggregate of the bank forming an 'optimal' output. Unfortunately, this method, along with others, suffers from the fact that only a small portion of a potential modeling space can be considered in any given realization of a filter bank. Additionally, the Magill filter experiences numerical underflow problems for long spans of data [2], [3]. Early attempts to utilize Magill filter banks for interplanetary orbit determination were reported by Burkhart and Bishop [4]. Another classical technique accommodates modeling errors by matching process noise and measurement noise statistics to the received data [5]. A recent application of this approach by Powell [6] successfully utilizes a simplex method, and thus does not require gradient information about the filter's dynamic and/or measurement models. Noise matching methods maintain tracking stability of the filter, however they do not address the fundamental issue of adjusting internal modeling assumptions that may have become suboptimal. Many other model optimization techniques exist, such as the recursive quadratic programming (RQP) approach investigated by Chaer, Bishop, and Ghosh [3],[7], however they are typically based on the existence of gradient information. These do not exist when considering model changes between discrete options (i.e., such as changing a flat plate to a cylinder). Adaptation using genetic algorithms (GA) is ideally suited for situations where the design space is complex and consists of mixed variable (discrete and continuous) because gradient information is not required. However, currently, they are not well suited for real time processing because their convergence is evolutionary in nature. Nevertheless, interplanetary navigation (especially during cruise) is typically a process that operates on spans of data that are days to weeks in length, and the filter tuning process (via analysis by a navigation team) can take days to weeks, as well. Thus, use of a genetic algorithm to assist in this process is warranted. Previously, Chaer, Bishop, and Ghosh [3], [7], and Chaer and Bishop [8] employed GAs for adaptive orbit determination during interplanetary cruise, however, their efforts focused on adjusting internal parameters (e.g. measurement and process noise) within individual filters of fixed structure. Using gating networks to regulate the filter bank, the GA operated in an outer loop with the performance of the individual filters represented by the gating network weights. In this new application, the GA is used to adapt the individual filter structure by effectively updating the spacecraft model itself. The filter bank is utilized as the GA population. This is a significantly different application of the GA from previously reported investigations. Also, for the first time, actual Deep Space Network (DSN) tracking data is used in the investigations, whereas previous studies relied on simulated DSN tracking data.

The current study employs a GA coupled with an extended Kalman filter (EKF) for model optimization. In particular, the solar radiation pressure (SRP) model of the Mars Pathfinder spacecraft is investigated using a 3 month span of tracking data during the cruise phase of the mission. The results obtained in this study are compared to the best model obtained by the MPF navigation team. During operations the issue of appropriate spacecraft component models was of a central concern. Shadowing of the backshell, coupled with a limited component selection complicated the team's search for an appropriate model. Given this experience, the design space selected for the GA search includes component selection of the backshell between either cylinders or flat plates. The size and orientation of these components is selectable. The GA design space also includes as selectable parameters the apriori covariances for the SRP component area scale factors. The GA operates on a population of individuals that are selected (initially at random) from this design space. Each individual processes the tracking data set via the EKF. The basis for the GA's fitness function is a normalized sample statistic of the output innovation sequence for each individual. Using the fitness values computed for each individual, the GA selects the parent population via a tournament method. For crossover, several strategies are investigated to determine the best method for quick convergence of the GA to a near optimal solution. The selected design space includes $1.44\text{E}+17$ distinct SRP models, and the results show that the GA is able to determine an SRP model with a fitness value that is $\sim 6\%$ better than the model selected by the MPF navigation team, and produces predicted residuals that are more stable.

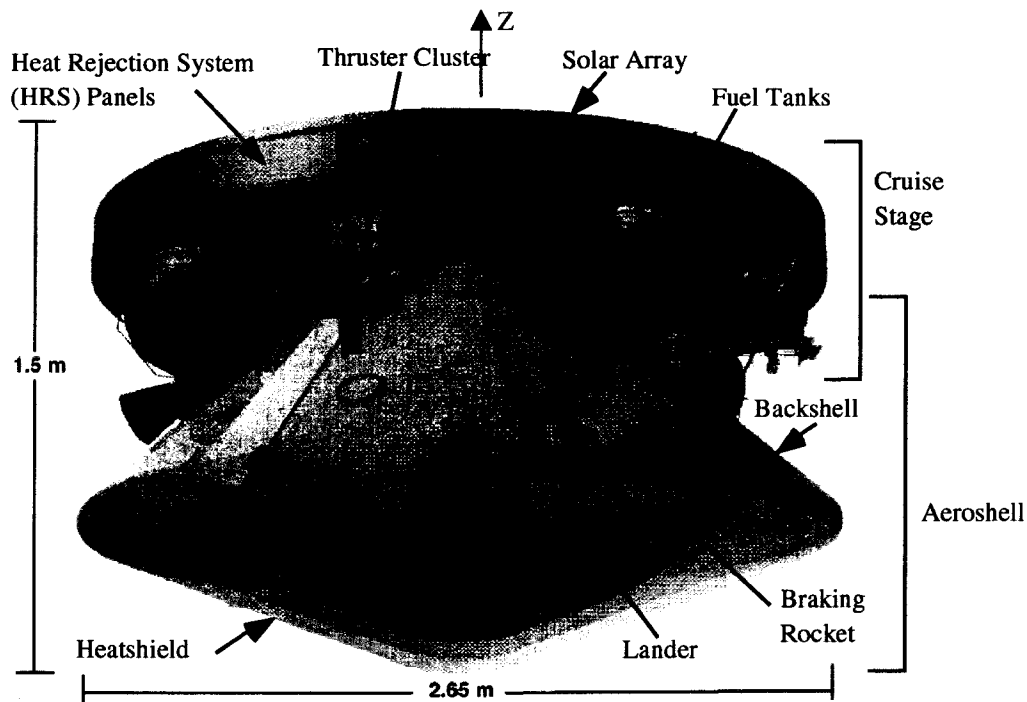


Figure 2: Mars Pathfinder spacecraft in cruise configuration [1].

MARS PATHFINDER NAVIGATION

The Mars Pathfinder spacecraft successfully landed on Mars on July 4, 1997 after launching 7 months earlier on December 4, 1996. The spacecraft trajectory was a Type 1 transfer from Earth to Mars and is shown in Figure 1. The passage of time is indicated on the trajectory via time ticks at 15 day intervals. The cruise phase trajectory correction maneuvers (TCMs) are indicated as well. Of particular interest in this study is modeling of the spacecraft dynamics during the cruise phase between TCM 2 on February 3, 1997 and TCM 3 on May 7, 1997, a span of 91 days. This span of data was selected for analysis because modeling issues regarding radiometric data quality, a significant effect during other data spans, was not a factor during this period. The configuration of the MPF spacecraft during cruise is shown in Figure 2. The MPF lander and its science payload were enclosed in an atmospheric deceleration module (backshell and heatshield) that was attached to a cruise stage. The cruise stage, mounted at the apex of the backshell, consists of a circular solar array panel, propellant and attitude control system hardware. Telecommunications and navigation tracking were performed using an X-band radio system and medium gain antenna, also on the cruise stage. Two way, coherent Doppler and ranging were the primary navigation data types utilized for trajectory determination.

A significant issue that the navigation team dealt with was the selection of an appropriate set of spacecraft components for solar radiation pressure modeling [1]. The software used to process the tracking data had modeling capabilities consisting of only standard shapes, such as flat plates and cylinders. However, the backshell's conical shape did not conform well to this set. Another factor adding to modeling difficulties was the fact that the backshell was partially shadowed by the solar array during most of the flight. Figure 3 shows the Sun-spacecraft-Earth angle γ during cruise. The spacecraft body fixed axes points the Z-axis outward and normal to the solar array. Since the spacecraft attitude is nominally pointed with the fixed Z-axis towards Earth, as γ becomes smaller the Sun illuminates more of the solar array, and the shadow cast by the array onto the backshell becomes larger. Examination of Figure 3 reveals that the backshell shadowing is a minimum at the start of the mission. It then increases to a maximum after ~ 70 days (~ 8 days into the period being investigated in this research), and then decreases after that. Over the

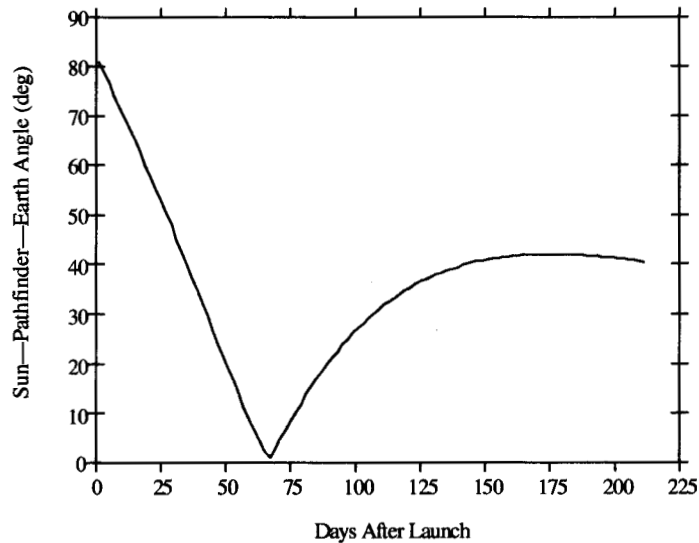


Figure 3: Sun-spacecraft-Earth angle γ for Mars Pathfinder cruise

course of the mission the navigation team experimented with many component models that tried to account for the impact of these geometric considerations sufficiently.

The process that the MPF navigation team utilized for model selection centered around filtering with all available data, then taking a solution based on the selected model and 'passing it through' the next few days of data. The model yielding the most stable pass through residuals would then be selected. During this time period the MPF navigation team arrived at a SRP model with components as indicated in Table 1. Of particular note is the change in the backshell model from a cylinder prior to April 16, 1997, to a flat plate after that. This illustrates a significant feature with the time dependent nature of SRP modeling resulting from the changing shadowing environment. It should be noted that later in the mission the team had to change the backshell model to be just a flat plate for the entire cruise because the solution using the model in Table 1 became extremely sensitive to the inclusion of a specific pass of data collected a week prior to Mars atmosphere entry.

Table 1: MPF navigation team optimal SRP model between TCM2 and TCM3

Spacecraft Part	Component Type fp = flat plate cyl = cylinder	Size	Active Span	Nominal Values for OD Filter Parameters		
				Specular Coefficient	Diffuse Coefficient	Area Scale Factor
Solar array	fp	Area 4.5 m ²	Entire Cruise	0.200	0.0208	1.063
Launch vehicle adapter	fp	Area 0.8 m ²	Entire Cruise	0.0	0.3001	1.067
HRS panels	cyl	0.273 m length, 1.35 m radius	Entire Cruise	0.0	0.3002	1.240
Backshell Model 1	cyl	2.3 m length, 1.35 m radius	Before 4/16/97	0.0	0.2989	0.4129
Backshell Model 2	fp	Area 2.7 m ²	After 4/16/97	0.0	0.3115	0.4358

GENETIC ALGORITHMS AND ADAPTIVE NAVIGATION

The experience of the MPF navigation team in finding an adequate SRP model illustrates the complexity of the navigation model design space and the ad hoc nature currently utilized to obtain an adequate model. It should be emphasized that the nature of the team's design decisions focused not only on filter parameter selection and their associated values, but also on selecting the underlying dynamic models (a structural change to the filter). Past approaches to adaptive filtering have typically focused on filter tuning by considering only modifications to model parameters values and/or a priori uncertainties. Structural changes to filters via modifications to underlying dynamic models, measurement models, and/or filter state vector components have been outside of the scope of most adaptive techniques. This is partially attributable to the fact that many design optimization methods are gradient based, and, thus, do not support discrete structural changes. An exception to this is the Magill filter bank, however, as noted previously, this method suffers from numerical underflow problems and a practical implementation of a bank is limited to a small set of all possible discrete filters. A genetic algorithm does not require gradient information nor is it limited by a small sample of the desired design space, hence it is a natural choice as a model optimization method.

A genetic algorithm is a computational representation of natural selection that bases its search and optimization on the analogy that an individual that is more *fit* to its environment is closer to an optimal design. In applying this analogy to the adaptive navigation problem an *individual* represents a specific filter realization of dynamic models, measurement models, state vector components, a priori uncertainty values, and associated noise process models. The particular variables and their associated range of values that have been isolated for analysis represent the design space. An example design space relevant to this study is spacecraft component selection for modeling the backshell. Each individual design has its variable values encoded into a representation, typically a binary string, that corresponds to its *chromosome*. The GA examines a *population* of individuals by analyzing the *fitness* of each individual, a metric that measures a selected figure of merit. In the case of adaptive navigation, the fitness value is based on the sample statistics of the residual sequence. With this information, the GA iterates on the members of the population from one *generation* to the next with the aim of improving the overall fitness of the population, where *improvement* is defined in this problem as *minimizing* the fitness value. The GA accomplishes this objective using three primary operators,

1. *Selecting* the more fit individuals of the population to become *parents*,
2. *Mating* the parents via a *crossover* operation that exchanges chromosome information to produce *children*,
3. *Modifying* the children's chromosomes via a low probability of *mutation* to ensure diversity of the population.

The children of the current generation then become the population for the next generation, and the process iterates until convergence or it is stopped. Furthermore to ensure the GA starts with a diverse image of the potential design space the initial population is seeded with a random set of individuals. This combination of deterministic rules for selection and combination coupled with probabilistic sampling creates a robust approach for arriving at the global minimum. Furthermore, the probabilistic elements of the GA help to prevent convergence towards local extrema.

THE ADAPTIVE NAVIGATION FITNESS FUNCTION

A general method for analyzing the fidelity of a filter that is processing measurement data is in the quality of its post fit residual sequence. Regardless of the filtering method (i.e., batch least squares, linearized or extended Kalman filter, etc.), a properly operating filter should produce a post fit residual sequence that is a zero mean white noise process [5]. For the application considered in this study an Extended Kalman (EKF) was utilized to process the 2 Way Doppler data collected by the Deep Space Network (DSN). Using the standard EKF algorithm, the residual value $r(\cdot)$ at time t_i is formulated as,

$$r(t_i) = z(t_i) - h(\hat{\mathbf{x}}(t_i^-), t_i), \quad (1)$$

where $z(\cdot)$ is a scalar measurement (in the present case a 2 Way Doppler value) taken at time t_i , $h(\cdot, \cdot)$ is the nonlinear measurement model, and $\hat{\mathbf{x}}(t_i^-)$ is the nonlinear state vector estimate propagated from the time of the prior measurement t_{i-1}^+ to the current time t_i^- (where the '-' indicates that the filter update with the current measurement has yet to occur). During mission operations, the measurement data is taken in passes when the satellite is in view of a DSN station. The time interval between measurements is typically uniform during a pass (however data editing can eliminate individual points within the pass), but the length of a pass and the interval between passes is typically not uniform because of changing geometry and operational schedules. The covariance associated with the residual in Eq. (1) takes the form,

$$w(t_i) = \mathbf{H}[t_i, \hat{\mathbf{x}}(t_i^-)] \mathbf{P}(t_i^-) \mathbf{H}^T[t_i, \hat{\mathbf{x}}(t_i^-)] + \sigma_z^2, \quad (2)$$

where $\mathbf{H}[t_i, \hat{\mathbf{x}}(t_i^-)]$ is the linearized measurement matrix and is defined as,

$$\mathbf{H}[t_i, \hat{\mathbf{x}}(t_i^-)] \equiv \left. \frac{\partial h[\mathbf{x}, t_i]}{\partial \mathbf{x}} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)}, \quad (3)$$

the predicted state covariance matrix is $\mathbf{P}(t_i^-)$, and σ_z^2 is the associated measurement noise variance (assumed to be time independent). A sequence of residual values $\{r(t_i), i=1, \dots, N\}$ that results from operating an EKF over a selected span of data can be used to formulate a normalized sample statistic ϕ that is a measure of the filter's performance. The statistic is computed using,

$$\phi = \frac{1}{N} \sum_{i=1}^N \frac{r(t_i)^2}{w(t_i)}. \quad (4)$$

It can be easily shown that the value that the statistic computed using Eq. (4) is related to the sum of the residual sequence's mean value squared and its associated variance, that is,

$$\phi = M^2 + V = \left(\frac{1}{N} \sum_{i=1}^N \frac{r(t_i)}{\sqrt{w(t_i)}} \right)^2 + \frac{1}{N} \sum_{i=1}^N \left(\frac{r(t_i)}{\sqrt{w(t_i)}} - M \right)^2, \quad (5)$$

where M represents the mean and V is the variance. Ideally an EKF operating nominally as compared to an EKF with a suboptimal model produces a residual sequence with a smaller mean value and variance. Thus, from Eq. (5), better filter models yield smaller values for ϕ . Finding a *global* optimal model $\hat{\mathbf{a}}$ for the EKF can now be posed in the form of a design objective using the statistic defined above:

$$\text{Find } \{\hat{\mathbf{a}} \in A \mid \phi_{\min} \equiv \phi(\hat{\mathbf{a}}) \leq \phi(\mathbf{a}), \forall \mathbf{a} \in A\}, \quad (6)$$

where \mathbf{a} is the vector of design parameters representing an individual that is a member of the selected design space A . Typically, a GA's fitness function for single objective optimization is defined as the sum of an objective (in this case the statistic ϕ) and an imposed penalty p , that is

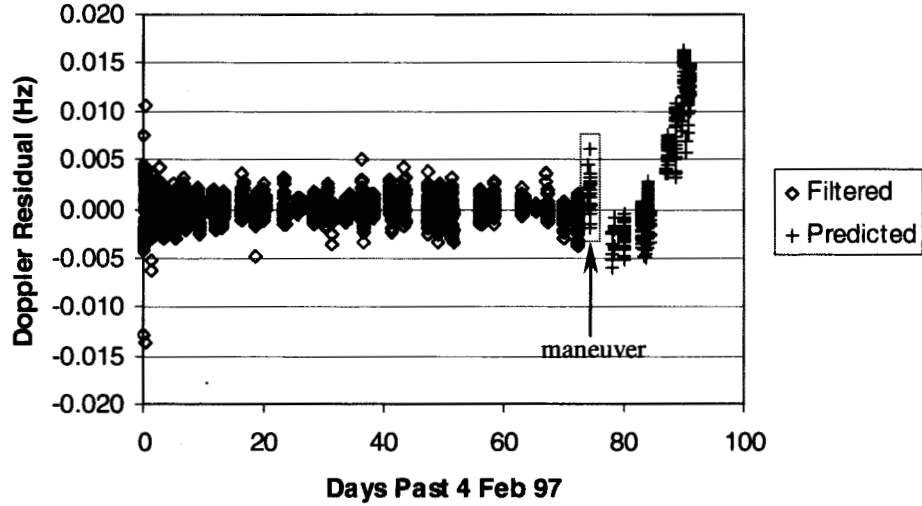


Figure 4: Doppler residual sequence using the MPF team's best SRP model ($f_{\text{MPF}} = 0.1470$).

$$f \equiv \phi + p \quad (7)$$

where the penalty function's definition is based on any imposed constraints. Since the problem that is being considered in this study does not have any constraints, the fitness function reduces to $f \equiv \phi$, and the design problem in Eq. (6) is equivalent to finding $\hat{\alpha}$ that minimizes the fitness f .

In the present problem the individual design selected by the MPF team listed in Table 1 is an example that yields a specific value for the fitness f . The data sequence consists of 2017 2-Way Doppler values collected between the DSN and the MPF spacecraft, and spans 91 days past February 4, 1997. This period lies between TCM 2 and TCM 3 on the trajectory. Figure 4 illustrates the residual sequence obtained on this data span using the SRP model of Table 1. Details of other filter model parameters can be found in Ref. [1]. Applying Eq. (4) to the first 1612 points of the sequence yields a fitness value of $f_{\text{MPF}} = 0.1470$. It is this result that the GA will try to improve upon as it searches through the design space. Also, note that not all 2017 data points were used to formulate the fitness. This is because only the first 72.5 days of data (1612 points) were filtered using the EKF. The EKF solution at the end of this span is then used to compute predicted measurements from days 74 to 91. These are differenced with the observed data collected during the same period to produce the predicted residual sequence shown in Figure 4. (Note that the pass of residuals on day 74 appears slightly biased due to the presence of a small correction maneuver that took place during the middle of the pass.) A stable solution should produce predictions that appear zero mean and white. It is clear from Figure 4 that this is not the case, there exists a distinctive divergence away from a zero mean. Ideally, the optimized model obtained by the GA will yield residuals that have not only a smaller value for f , but also have a filter solution that produces more stable predicted residuals, as well.

THE MPF SOLAR RADIATION PRESSURE DESIGN SPACE

Perhaps the most crucial part of using a GA for model optimization is the selection of an appropriate design space and the formulation of the parameter vector α . An immediate practical concern is selection of a minimal set of design parameters that adequately characterize the design problem under investigation. The number of design parameters and the associated bit level resolution necessarily increase the length of an individual chromosome, and subsequently increase the size of the design space. For example a chromosome length of 57 bits yields a design space of $2^{57} \approx 1.441 \times 10^{17}$ individual designs. The required

Table 2: The SRP design space for the MPF spacecraft configuration for modeling solar radiation pressure.

Spacecraft Part	Type: fp	fp parameter: $\theta \in (30^\circ, 65^\circ)$	fp parameter: $A \in (1 \text{ m}^2, 3 \text{ m}^2)$	Active Span (T_i, T_{i+1}) \in ($T_i, T_i + 64$ days) Note: $T_0 = 4$ Feb 97	Area Scale Factor Sigma $\sigma_{SF} \in (0, 2)$
	Type: cyl	cyl parameter: $r \in (.5 \text{ m}, 1.4 \text{ m})$	cyl parameter $h \in (.5 \text{ m}, 2.5 \text{ m})$		
Solar array	fp	$\theta = 0^\circ$	$A = 5.2 \text{ m}^2$	Entire Cruise	3 bits \Rightarrow $\Delta\sigma_{SF} = .29$
Launch vehicle adapter	fp	$\theta = 0^\circ$	$A = .8 \text{ m}^2$	Entire Cruise	3 bits \Rightarrow $\Delta\sigma_{SF} = .29$
HRS panels	cyl	$r = 1.35 \text{ m}$	$h = .273 \text{ m}$	Entire Cruise	3 bits \Rightarrow $\Delta\sigma_{SF} = .29$
Backshell Model 1 Active (T_0, T_1)	1 bit	3 bits \Rightarrow $\Delta\theta = 5^\circ, \Delta r = .13 \text{ m}$	3 bits \Rightarrow $\Delta A, \Delta h = .29 \text{ m}^2, \text{ m}$	6 bits \Rightarrow $\Delta T = 1 \text{ day}$	3 bits \Rightarrow $\Delta\sigma_{SF} = .29$
Backshell Model 2 Active (T_1, T_2)	1 bit	3 bits \Rightarrow $\Delta\theta = 5^\circ, \Delta r = .13 \text{ m}$	3 bits \Rightarrow $\Delta A, \Delta h = .29 \text{ m}^2, \text{ m}$	6 bits \Rightarrow $\Delta T = 1 \text{ day}$	3 bits \Rightarrow $\Delta\sigma_{SF} = .29$
Backshell Model 3 Active (T_2, T_3)	1 bit	3 bits \Rightarrow $\Delta\theta = 5^\circ, \Delta r = .13 \text{ m}$	3 bits \Rightarrow $\Delta A, \Delta h = .29 \text{ m}^2, \text{ m}$	6 bits \Rightarrow $\Delta T = 1 \text{ day}$	3 bits \Rightarrow $\Delta\sigma_{SF} = .29$

population size increases as the chromosome grows longer, and the number of generations required for convergence increases as the design space becomes larger. Hence, given the evolutionary nature of the GA, the time to convergence can be a significant practical concern for long chromosomes. In the current problem a given filter realization processing the given data set on an Hewlett Packard 360 workstation took anywhere from 2.5 ~ 4.1 minutes of dedicated CPU time, the result is a GA run that can take many days to a week before completing. Hence, to minimize this time, a minimal set of design parameters is sought. Table 2 lists the design parameters selected for analysis, the portion that is highlighted in gray represent the specific variables to be encoded into a binary string to form the chromosomes. The areas not highlighted are parameters that remain fixed for each individual (in addition to other filter parameters not listed). Associated with each design variable is the range of values that it can take and the number of bits (genes) that are used to encode it into a binary string. Using the combination of the range ($\alpha_{\min}, \alpha_{\max}$) and the number of bits l specifies the resolution of the parameter $\Delta\alpha$, that is,

$$\Delta\alpha = \frac{\alpha_{\max} - \alpha_{\min}}{2^l - 1}. \quad (8)$$

The particular MPF SRP model parameters selected include a backshell model that can change with time as shadowing characteristics change. Three consecutive transitions are allowed, and each can be active for a span of time from 0 to 64 days. Since these add sequentially, it is possible for the backshell models to last for a total period that is anywhere from 0 days long up to 192 days (119.5 days longer than the span of tracking data). It is up to the GA to determine which combination of time spans best models the backshell. A novel feature of the design space centers on a concept pioneered by Crossley [9] that associates strings of genes together according to a selected design feature. Crossley called this gene association *sex-limited inheritance*. Specifically, the backshell model type can be either a flat plate or a cylinder, and, depending on the particular selection that is made, the subsequent two variables are an angle and area for the flat plate, or a radius and height for the cylinder. Note that θ is the angle between the flat plate outward normal vector and the spacecraft Z axis (nominally Earth pointed), with the vector lying in the YZ spacecraft plane (nominally transverse to the orbit plane). In the current application it is possible to formulate a selection between multiple type dependencies by allowing the type variable to have more than a 1 bit representation (even though only two types are specifically investigated), as such, the term *type-limited inheritance* will be used rather than sex-limited inheritance. Finally, the a priori uncertainty associated with the area scale

Table 3: MPF SRP chromosome structure.

SA σ_{SF}	LVA σ_{SF}	HRS σ_{SF}	BS1 type	BS1 1 st para.	BS1 2 nd para.	BS1 span	BS1 σ_{SF}	BS2 type
3 bits	3 bits	3 bits	1 bit	3 bits	3 bits	6 bits	3 bits	1 bit

BS2 1 st para.	BS2 2 nd para.	BS2 span	BS2 σ_{SF}	BS3 type	BS3 1 st para.	BS3 2 nd para.	BS3 span	BS3 σ_{SF}
3 bits	3 bits	6 bits	3 bits	1 bit	3 bits	3 bits	6 bits	3 bits

factors for all the components are selected as design parameters (the scale factors for each component are nominally included in the filter state vector). All other filter parameters remain fixed and are the same as the filter parameters used during operations. The selected design variables and number of genes associated with each variable yield a chromosome with a length of 57 bits (genes).

A key consideration with formulating the structure of the chromosome is in ordering these genes. The strength of a GA's search versus a pure stochastic search lies with its ability to retain knowledge of desirable features from one generation to the next. This is the concept of schema processing. A schema can be defined as a gene pattern describing a common feature in a subset of the population. For instance, if a backshell 2 model of a flat plate with an angle of 45° consistently produces better fitness values over other models, then a gene pattern representing this selection can become prevalent in a number of the population via the GA selection operation. In this fashion, the GA retains desirable features from one generation to the next. It has been shown analytically, using the fundamental theorem of genetic algorithms [10], that schema of low order and length (i.e., the number of distinguishing genes and the distance between the last and first gene, respectively) assist in the convergence process. Hence, the GA designer should take care to put like features near each other in the construction of the chromosome string. For example, all genes associated with a single backshell model should be next to each other. With this in mind, the chromosome structure selected for this study takes the form indicated in Table 3.

A final aspect regarding the construction of the chromosome is in the choice of encoding. A simple binary encoding is a classic choice, however this method presents a difficulty in the form of Hamming Cliffs [10]. Hamming Cliffs are associated with the number of bit flips required when incrementing an integer by one. As an example, consider the binary representation of 3, which is 011, and 4, which is 100. Incrementing the integer value by one changes all three of the bit positions. Thus the GA, which operates on individual bits in its crossover operation, will have difficulty in incrementing this example from 3 to 4 since three bit operations are required. The result is a Hamming Cliff. Furthermore, if the fitness levels change significantly between these two representations, the GA will have difficulty in finding this change. To circumvent this a Gray code is selected for encoding the binary strings. The salient feature of a Gray code is consecutive integers result in binary representations that change only one bit. In the prior example, the Gray code representation of 3 is 010 and for 4 it is 110, only the left most bit has changed. Details of the Gray code algorithm can be found on Weisstein's web page [11].

THE GENETIC ALGORITHM OPERATORS

Three main operators form the basis for a GA algorithm search – selection, crossover, and mutation. A binary, single branch tournament approach is used to select the parents in the current population for mating. Unlike other popular methods, such as roulette wheel selection, this approach is a competitive strategy that is amenable to minimization problems without fitness scaling. It is also easily adaptable to multiobjective design via adding another branch to the tournament scheme. For instance, the two branch scheme can be easily formulated to minimize an objective based on 2 Way Doppler data *and* another objective function based on range data (a topic under investigation in the continuation of this research). Parent selection proceeds as follows:

1. Place the entire population into a *pot*,

2. Select two individuals at random and without replacement,
3. Compare their fitness values, the individual with the lower fitness is placed into a *parent pool*,
4. Repeat this process until the pot is empty – the parent pool is now half full,
5. The pot is refilled with the entire current population.
6. Repeat steps 2 through 4. This fills the parent pool completely and readies the GA for the crossover operation.

Note that this sequence of operations guarantees that the overall fitness of the parent pool is better than the population.

After selection, individuals in the parent pool are randomly selected without replacement, two at a time, and *mated* in the crossover operation to form two children. Again there are numerous crossover approaches to choose from, in the present study, results generated using single point crossover and uniform crossover are examined. In single point crossover, the chromosomes of two parents are selected and if a random number draw between (0,1) is less a threshold value (set to 0.9 in this study) they exchange genetic material. If there is no exchange the parents become children (they are *clones*). If there is an exchange, it takes place at a random location along the parents chromosome string. The genetic material before this location remains unchanged, and the genetic material after it is switched. The result is two children that inherit genetic material from both parents. For example, if the random location along the chromosome is 23, then the parents' genes from 1 to 23 remain fixed, and the genes from 24 to 57 are swapped. The other strategy that is investigated is uniform crossover. In this approach each child string receives each gene from the first parent or the second with a 50% chance. For example, if a fair coin toss came up "heads" the first child would inherit its first gene from the first parent and the second child would inherit its first gene from the second parent. If the toss were "tails", the first child would inherit its first gene from the second parent and the second child would inherit its first gene from the first parent. Each gene location is examined for crossover in this manner, and the resulting children will inherit traits of both parents through the binary chromosomes. Empirical evidence suggests that single point crossover is better at retaining long-length and long order schema than uniform crossover (exploiting information known to be good). However, uniform crossover is better at exploring the design space to find more fit designs. [12]

The mutation operation occurs with a very low probability according to the following rule,

$$P_m = \left| \frac{1}{N_{POP}} - \frac{1}{N_{GENES}} \right| / 2 + \min \left(\frac{1}{N_{POP}}, \frac{1}{N_{GENES}} \right), \quad (9)$$

where N_{POP} is the size of the population and N_{GENES} is the number of genes in a chromosome. In the present study N_{GENES} is equal to 57 and N_{POP} is set to 172 (which is $\sim 3 \times N_{GENES}$), thus the probability of mutation becomes 0.0117. After the child strings have been formed, a mutation may occur that will change a gene to its binary opposite. Each gene of a child's chromosome is subject to mutation. Mutation ensures diversity because it can allow for a new trait to be introduced that was not present in a child's parents. It also assists the GA from converging prematurely into a local minimum. Finally, the GA utilizes elitism. Since a best individual (as determined by the selection operation) can be mutated, it is possible to lose desirable genetic information. To prevent this from happening, the best individual in a generation is retained, and, after mutation, put back into the population by replacing a randomly selected individual. The outcome of the selection, crossover, and mutation operations is a new population ready for the next generation of GA processing.

Now, with the selection operator favoring good individuals, the crossover operator combining features of good individuals, and the mutation operator ensuring diversity the GA population moves towards the globally optimal design. The GA search continues until convergence, where this is defined in the present context as the fitness function difference does not change by a defined percentage for a selected number of generation (1/1000% for 10 generations in the current implementation).

Table 4: Comparison of best solutions.

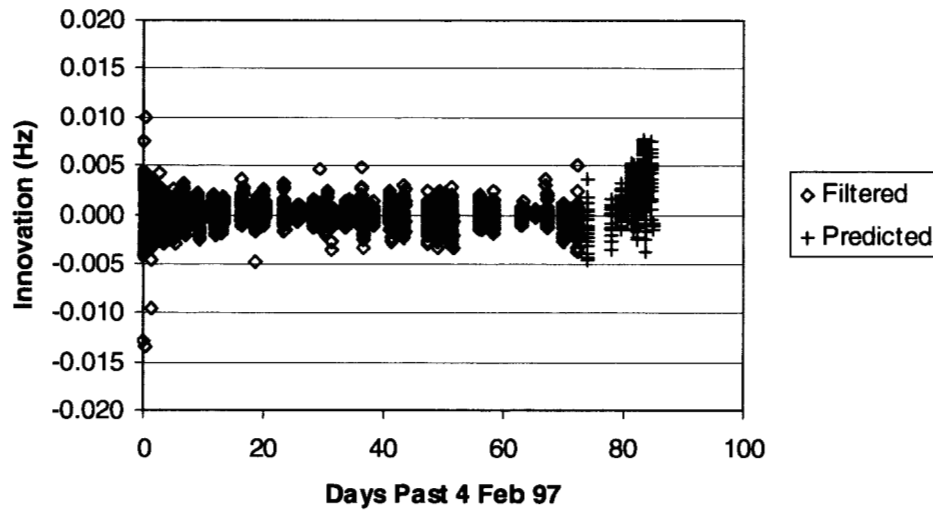
Solution	SA σ_{SF}	LVA σ_{SF}	HRS σ_{SF}	BS1 type	BS1 1 st para.	BS1 2 nd para.	BS1 start	BS1 σ_{SF}	BS2 type
MPF Nav Team Best SRP	.4	.4	.4	cyl	$r = 1.35$	$h = 2.3$	Epoch	.4	fp
GA SRP: Uniform Cross.	2.00	2.00	1.15	cyl	$r = .5$	$h = 1.36$	Epoch	1.43	cyl
GA SRP: Single Cross.	1.72	2.00	.579	cyl	$r = .76$	$h = 1.36$	Epoch	1.15	cyl

Solution (cont'd)	BS2 1 st para.	BS2 2 nd para.	BS2 start	BS2 σ_{SF}	BS3 type	BS3 1 st para.	BS3 2 nd para.	BS3 start	BS3 σ_{SF}
MPF Nav Team Best SRP	$a = 45^\circ$	$A = 2.7$	4/16/97	.4	Not Applicable				
GA SRP: Uniform Cross.	$r = .76$	$h = 2.5$	3/4/97	1.15	cyl	$r = 1.27$	$h = 2.21$	3/30/97	1.43
GA SRP: Single Cross.	$r = .89$	$h = 2.5$	2/23/97	2.00	cyl	$r = .63$	$h = 2.21$	4/11/97	1.72

RESULTS

The results of the genetic algorithm search over the design space identified in Table 2 using the chromosome structure defined in Table 3, and the operators discussed in the prior section produce an SRP model that is better than the MPF team result in several contexts. Two cases are considered, a GA run using uniform crossover and a second using single point crossover. In each case, convergence was achieved in relatively few generations, 12 with 10 additional generations for verification of convergence. Table 4 lists the solution details, and includes the MPF teams result for comparison. The associated filter residuals (and associated fitness values) resulting from these models are shown in Figure 5. Finally, the convergence characteristics of these two GA runs are illustrated in Figure 6. There are a number of features of both GA solutions that are immediately apparent. The backshell configurations selected by the GAs are consistently cylinders, as opposed to a flat plate for one of the MPF navigation team selections. Consider that the backshell, in actuality, has a conical shape, thus a cylinder would seem to be a better representation of its geometry as opposed to a flat plate. This is indeed the case as determined by the GA. Furthermore, the later backshell models (2 & 3) are larger than the first backshell component. Recall, from Figure 3 that during the later part of cruise the Sun-spacecraft-Earth angle grows, and, hence, the backshell is shadowed less by the solar array. Given this, it is reasonable to expect a cylinder with more surface area would lead to a better representation of the backshell later in the run, again, the GA arrives at this conclusion. Therefore, these GA solutions appear more intuitive from physical point of view. This conclusion is also supported quantitatively, the fitness values for both GA solutions ($f_{uni} = 0.1387$ and $f_{single} = 0.1389$) are nearly 6% better than the MPF team result of $f_{MPF} = 0.1470$. However, perhaps more significant, is both GA solutions produce predicted residuals that are not divergent (i.e., more stable) as compared to their MPF counterpart. A stable result is a strong indicator that the underlying models are more representative of reality, because the predicted measurements based on these models continue to match the real tracking data quite well for several weeks. Given the apparent divergence of the MPF team's result in Figure 4, it is not surprising that later in the mission this model had to be adjusted to be less sensitive to specific data. Finally, examination of Figure 6 shows that the zeroth generation (which is populated randomly) immediately produces an individual that is better fit than the MPF solution. The next 12 generations iterate on this result to fine-tune the solution. These fortuitous zeroth generation results stem, in part, by the selection of a fairly narrow design space. However, the design space selection was guided by realizing that any feasible solution should start in a region (of all possible designs) near a physically reasonable conception of the spacecraft components. Thus, it is not surprising to find good individuals quickly. The strength of the GA in this search was in fine tuning the selected SRP, with the

a.) Uniform Crossover ($f_{uni} = 0.1387$)



b.) Single Point Crossover ($f_{single} = 0.1389$)

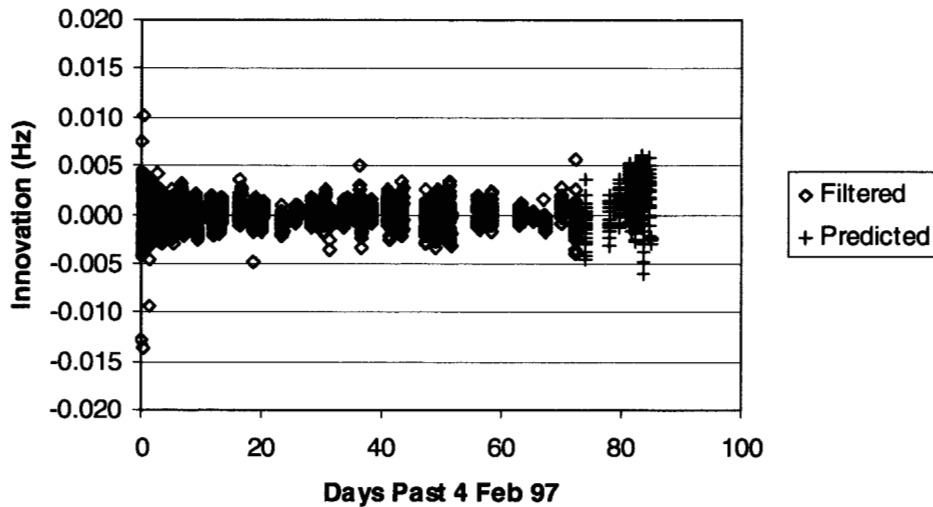


Figure 5: Residuals resulting from best GA model selected using different crossover strategies

crucial result that the final solution produces a more stable predicted residual sequence than the models arrived at using a manpower intensive, ad hoc search.

CONCLUSIONS

Adaptive navigation using genetic algorithms offers a promising solution to the problem of optimal filter model selection. The results of this study have illustrated that a systematic search through a predefined design space is capable of finding better filter models than the traditional manpower intensive approach. The GA is capable of making a myriad of design decisions that would be prohibitive if those decisions were made by a human. The result is a powerful tool to assist an analyst in determining an optimal filter model. Of course, convergence of the GA is predicated on a thoughtful construction of the GA design space, chromosome structure, and selection of the GA operators and their associated parameters.

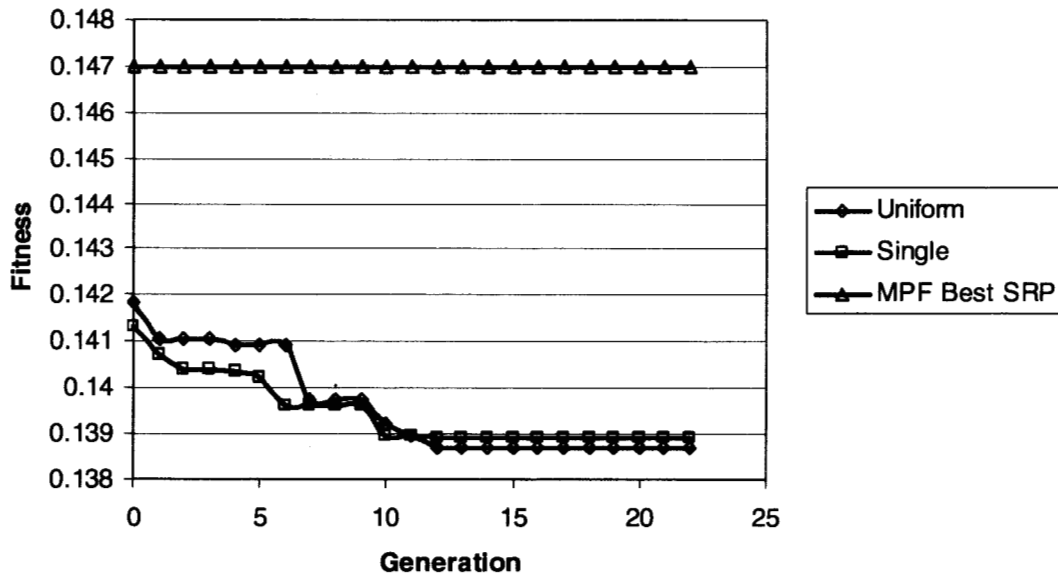


Figure 6: Convergence of GA runs to best solutions

This aspect of the GA based adaptive approach is well served by a careful analysis of the physical characteristics of the problem to be solved.

ACKNOWLEDGEMENTS

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The work of the second and third authors was supported by CIT/JPL grant 959577. The authors wish to thank Dr. L. Alberto Cangahuala, Dr. William Crossley, and Mr. Ted Drain for their helpful input and discussions regarding this work.

REFERENCES

1. Kallemeyn, P. D., Vaughn R. M., Spencer D. A., Braun, R. D., "Mars Pathfinder Navigation Report," Jet Propulsion Interoffice Memorandum, 312.A/98-030 (Internal Document), January 9, 1998.
2. Magill, D. T., "Optimal Adaptive Estimation of Sampled Stochastic Processes," *IEEE Transactions on Automatic Control*, Vol. AC-10, No. 4, Oct. 1965, pp. 434 – 439.
3. Chaer, W., Bishop, R. H., and Ghosh J., "Hierarchical Adaptive Kalman Filtering for Interplanetary Orbit Determination," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 34, No. 3, 1998, pp. 1-14.
4. Burkhart, P. D. and Bishop, R. H., "Adaptive Orbit Determination for Interplanetary Spacecraft," *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 19, No. 3, 1996, pp. 693-701.
5. Maybeck, P. S., *Stochastic Models, Estimation, and Control*, Vol. 1. Academic Press, 1979.
6. Powell, T. D., "Automated Tuning of an Extended Kalman Filter Using the Downhill Simplex Algorithm," Paper AAS 99-368, AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska, 16-19 August 1999.

8. Chaer, W. S. and Bishop, R. H., "Adaptive Kalman Filtering with Genetic Algorithms," *Advances in the Astronautical Sciences*, eds. R. J. Proulx et al., Vol. 89, pp. 141-155, 1995.
9. Crossley, W. A., "Using Genetic Algorithms as an Automated Methodology for Conceptual Design of Rotorcraft," Ph.D. Dissertation, Arizona State University, Tempe, AZ, Aug. 1995.
10. Goldberg, D. E., *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989.
11. Weisstein, E., *Eric Weisstein's World of Mathematics*, <http://mathworld.wolfram.com/GrayCode.html>.
12. Crossley, W. A., "Class notes AAE590G Multidisciplinary Design Optimization in Aerospace Engineering," Purdue University, Spring 1997.